

Effects of Adaptive Protective Behavior on the Dynamics of Sexually Transmitted Infections

Michael A. L. Hayashi^a, Marisa C. Eisenberg^{a,b}

^a*Department of Epidemiology, University of Michigan, Ann Arbor*

^b*Department of Mathematics, University of Michigan, Ann Arbor*

Abstract

Sexually transmitted infections (STIs) continue to present a complex and costly challenge to public health programs. The preferences and social dynamics of a population can have a large impact on the course of an outbreak as well as the effectiveness of interventions intended to influence individual behavior. In addition, individuals may alter their sexual behavior in response to the presence of STIs, creating a feedback loop between transmission and behavior. We investigate the consequences of modeling the interaction between STI transmission and prophylactic use with a model that links a Susceptible-Infectious-Susceptible (SIS) system to evolutionary game dynamics that determine the effective contact rate. The combined model framework allows us to address protective behavior by both infected and susceptible individuals. Feedback between behavioral adaptation and prevalence creates a wide range of dynamic behaviors in the combined model, including damped and sustained oscillations as well as bistability, depending on the behavioral parameters and disease growth rate. We found that disease extinction is possible for multiple regions where $\mathcal{R}_0 > 1$, due to behavior adaptation driving the epidemic downward, although conversely endemic prevalence for arbitrarily low \mathcal{R}_0 is also possible if contact rates are sufficiently high. We also tested how model misspecification might affect disease forecasting and estimation of the model parameters and \mathcal{R}_0 . We found that alternative models that neglect the behavioral feedback or only consider behavior adaptation by susceptible individuals can potentially yield misleading parameter estimates or omit significant features of the disease trajectory.

Keywords: transmission model, behavioral dynamics, game theory, sexually transmitted infections

Email addresses: mhayash@umich.edu (Michael A. L. Hayashi), marisae@umich.edu (Marisa C. Eisenberg)

1. Introduction

In spite of advances in treatment and prevention, sexually transmitted infections (STIs) remain endemic worldwide. The CDC estimates that 20 million new cases occur annually in the United States alone [1], incurring a total cost of \$16 billion for treatment and care. Globally, treatable STIs are responsible for approximately 500 million new cases per year [2], while an estimated 35 million individuals currently live with HIV. These statistics underscore the importance of understanding the dynamics that drive and sustain STI transmission. To this end, mathematical epidemiology has made substantial progress investigating the role of contact patterns such as age-structure, sexual networks, and levels of sexual activity [3, 4]. However, many open questions remain in understanding the feedback relationship between behavioral change and disease dynamics.

From a behavioral standpoint, sexually transmitted diseases are noteworthy as they require a direct and intimate interaction between individuals. As a result, many common preventative measures, such as condom use, are not determined unilaterally [5, 6]. In addition, assuming individuals form preferences over protective behaviors based on the associated costs and benefits, we would expect these behaviors to respond to the risk of infection as an outbreak progresses [7–9]. While changes in risky or protective behavior can amplify or suppress outbreaks, adoption of these behaviors can in turn be driven by the spread of disease, as demonstrated by increased testing and condom use among men who have sex with men in response to the HIV outbreak in the US [10, 11].

Methods from game theory provide a framework with which to capture this feedback, grounded in well-established mathematical and economic theory. The resulting economic-epidemiological models can explicitly represent the decision process of individuals either in direct interactions (e.g. sexual encounters) or population interactions (e.g. vaccination behavior) [12–18]. Including the effects of behavioral change on STI dynamics has been primarily motivated by the HIV epidemic among men who have sex with men (MSM) in the 1980s and 1990s, but this modeling approach is relevant for the study of other pathogens and communities as well. Indeed, in 2013 the WHO highlighted the need to study behavioral change in order to design effective interventions [2].

Many economic-epidemiological models rely on two behavioral assumptions that are worth consideration. The traditional game theoretic framework assumes that all actors are fully rational, responding

optimally at every stage of play [19]. While convenient, the rationality assumption remains a subject of debate in economic literature. Empirical studies note circumstances in which individual behavior appears to depart from a strict payoff maximization foundation [20–22]. In addition, most models (e.g. [23] and [24]) assume that only susceptible individuals make choices regarding protective behavior. This assumption is the result of representing the cost-benefit calculus of individuals as a tradeoff between various private costs of protective behavior and the risk of infection. However, infected individuals may have incentives to reduce contact or use protection as well, motivated by altruism, self interest, or other factors [25, 26]. This has been the focus of several intervention strategies in practice [27, 28].

In this paper, we propose a model of combined behavioral and disease transmission dynamics that uses the outcome of sexual interactions between susceptible and infected individuals to determine the effective contact rate for a mass action model of disease transmission. The combined model bears some similarity to the behavior-disease model proposed in [13]. There are, however, several critical distinctions. We use a deterministic ODE framework, while our game-theoretic model collapses the protection-use game to a single interaction instead of a multi-stage negotiation. In addition, similar to [14] and [17] we use evolutionary dynamics to represent the process of behavioral change over time, allowing for non-optimal but potentially more realistic behaviors. Unlike the inductive reasoning game developed by Breban et al. [29], our behavioral dynamics only explicitly consider the current state instead of a history of actions. While this approach loses some realistic features, it still allows us to relax the assumption of full rationality while also providing a convenient mathematical formulation for the combined model [30, 31].

2. Model

The Susceptible-Infectious-Susceptible (SIS) model has been studied extensively as a simplified representation of bacterial sexually transmitted diseases [32–34]. The model equations are

$$\begin{aligned}\dot{S} &= \gamma I - \beta SI, \\ \dot{I} &= \beta SI - \gamma I,\end{aligned}\tag{1}$$

where S is the fraction of the population which is susceptible, I is the fraction infected, β is the effective contact rate; the product of the rate of sexual partner acquisition and the probability of disease transmission from an infected to a susceptible partner, and γ is the rate of recovery or treatment. The basic reproductive ratio is

$$\mathcal{R}_0 = \frac{\beta}{\gamma}. \quad (2)$$

The disease free equilibrium (DFE) occurs if $\mathcal{R}_0 < 1$. Otherwise the endemic prevalence is

$$I^* = 1 - \frac{\gamma}{\beta}. \quad (3)$$

In order to capture the potential for individuals to adapt their protective behaviors over the course of an outbreak, we define a Bayesian game [35] between a pair of players. In Appendix A, we give a brief overview of the definitions and structure of Bayesian games, with more complete treatments given in [19, 35, 42]. The payoffs for the game depend on the disease states of both players, which are considered private information. Players must infer the type of their partner, reflecting realistic uncertainty about serostatus [26, 36–39]. Consistent with the notation for the SIS model, a player may be one of two types chosen from the type space $\Theta = \{S, I\}$. Each player chooses between using protection (P) or no protection (U) for a given sexual encounter. If both players select the same action, the outcome of the game is the same as the chosen action. We assume that if both players select different actions the encounter does not proceed and the effective contact rate for the pair of players is 0. In order to characterize the strategy space for this game, it is convenient to use the type-contingent notation $\sigma_j(\theta_i) \in \{P, U\}$ to denote the action player i would choose if she was of type θ_i . A complete strategy for a player then specifies a pair of actions $\sigma_j := \sigma_j(\theta_i = S)\sigma_j(\theta_i = I) \in \{PP, PU, UP, UU\}$. Without loss of generality, the type-dependent payoff (utility) to player 1 for a given pair of actions and types is written $u_1(\sigma_j(\theta_1), \sigma_k(\theta_2), \theta_1, \theta_2)$ (where the first entry specifies player 1's action assuming type θ_1 and the second entry player 2's action assuming type θ_2). Then player 1's overall expected payoff for a particular strategy profile (i.e. for a pair of type-contingent strategies σ_j and σ_k for each player) is the

		Player 2			
		P	U		
Player 1	P	b	c	Player 1	P
	U	c	a		U
		(a)			

		Player 2			
		P	U		
Player 1	P	b	c	Player 1	P
	U	c	d		U
		(b)			

Figure 1: The two variants of the protected sex game for the cases a: $\theta_i = \theta_j$ and b: $\theta_i \neq \theta_j$.

double expectation of $u_1(\sigma_j(\theta_1), \sigma_k(\theta_2), \theta_1, \theta_2)$ over both players potential types [19], that is,

$$E[u_1(\sigma_j(\theta_1), \sigma_k(\theta_2), \theta_1, \theta_2)] = \sum_{\theta_1} Pr(\theta_1) \left[\sum_{\theta_2} p_1(\theta_2|\theta_1) u_1(\sigma_j(\theta_1), \sigma_k(\theta_2), \theta_1, \theta_2) \right] \quad (4)$$

where $E[u_1(\sigma_j(\theta_1), \sigma_k(\theta_2), \theta_1, \theta_2)]$ is sometimes written more simply as $E[u_1(\sigma_j, \sigma_k)]$. The prior probabilities for each player being of either type, $Pr(\theta \in \{S, I\})$, are given by the distribution of susceptible and infected individuals in the population. We assume that partner selection is not assortative in disease type, so the belief for player 2 being type θ_2 given player 1 being type θ_1 is $p_1(\theta_2|\theta_1) = Pr(\theta_2)$.

Fig. 1 shows the two possible payoff matrices for the protected sex game that determine the type-dependent payoff terms $u_i(\sigma_j(\theta_1), \sigma_k(\theta_2), \theta_1, \theta_2)$ in Equation 4. For a concrete example, we compute the expected payoff to player 1 of playing strategy $\sigma_j = PU$ if player 2 picks $\sigma_k = UU$ when 30% of the population is infected and we choose $a = 1, b = 0.75, c = 0.5, d = 0.25$. From Equation 4 we have

$$\begin{aligned} E[u_1(PU, UU)] &= 0.7(0.7 \times 0.5 + 0.3 \times 0.25) \\ &\quad + 0.3(0.7 \times 0.25 + 0.3 \times 1) \\ &= 0.44. \end{aligned} \quad (5)$$

In general, we suppose the specific payoff entries satisfy $a > b > c > d$. Thus, individuals prefer unprotected sex to protected sex with a partner of the same disease status, preferring either to action pairs resulting in no sexual encounter. However, individuals prefer protected sex to all other outcomes when their partner is of a different disease type. This is intended to capture the notion that both susceptible and infected individuals have some incentive to avoid infection or transmission respectively. We explore the alternate case where infected individuals do not distinguish between susceptible and infected partners in Appendix E.

Consequently, for the protected sex game, the time-varying payoff matrix M is a 4×4 square matrix with elements given by Eq. (4). We can compute the Bayes-Nash equilibrium for any given choice of a, b, c, d, S, I . One case of particular interest is at the disease free equilibrium, $I^* = 0$. Here the game reduces to the 2×2 game with payoffs as in Figure 1a. This game has two pure strategy Nash equilibria, (P, P) and (U, U) . However, the (U, U) equilibrium in the reduced game is both payoff and risk dominant. Since the $\sigma(\theta_i = I)$ actions do not contribute to the expected payoff, any mixture of the type-contingent strategies UP and UU (similarly PP and PU) is a Nash equilibrium in the full model. We will revisit this point shortly.

Since we are ultimately interested in the dynamics of behavior at more than a single disease state, we use methods from evolutionary game theory to couple the dynamics of strategy change over time to the disease trajectory. This allows us to close the feedback loop between behavior and disease dynamics. In particular, we use replicator-mutator dynamics with a linear fitness function [31, 40]. The distribution of strategies in the population using the replicator-mutator approach depends both on the existing distribution of strategies and their expected payoffs, as well as a small degree of random strategy choice. This allows us to capture the notion that individuals may not respond immediately or strictly optimally to the presence of an outbreak, as payoff-suboptimal strategies may remain frequent in the population for some time. In addition, the possibility of mutation prevents any strategy from becoming extinct. In the behavioral context, this can be thought of as allowing for a small amount of random choice. This feature is particularly important as the replicator equation has stable steady states at strict pure strategy Nash equilibria, such that behavior cannot change after fixation on a particular Nash equilibrium even if disease conditions change. Both players in the protected sex game have the same action and type sets, so the relative frequency of the j th strategy in a large population is $f(\sigma_j, t)$. The replicator-mutator equation for a given strategy σ_j is

$$\dot{f}(\sigma_j) = s \left[\sum_k f(\sigma_k, t) q_{kj} \phi(\sigma_k, t) - \bar{\phi}(t) f(\sigma_j, t) \right], \quad (6)$$

where q_{kj} is the probability that an individual playing σ_j switches to σ_k , $\phi(\sigma_j, t) = \sum_k E[u(\sigma_j, \sigma_k)] f(\sigma_k, t)$ is the fitness of σ_j , $\bar{\phi}(t) = \sum_j \phi(\sigma_j, t) f(\sigma_j, t)$, and $s \in [0, \infty)$ is a scaling term that determines the

speed of behavior change. In vector-matrix form, this can be written as

$$\dot{\mathbf{f}} = s[D_{\mathbf{f}}Q(M\mathbf{f}) - (\mathbf{f}^T M\mathbf{f})\mathbf{f}], \quad (7)$$

where M is the 4×4 payoff matrix with elements $m_{jk} = E[u_1(\sigma_j, \sigma_k)]$, $\mathbf{f} = (f(PP, t), f(PU, t), f(UP, t), f(UU, t))$, Q is the mutation matrix and $D_{\mathbf{f}} = \text{diag}(\mathbf{f})$. Note that when $Q = I$, the evolutionary dynamics are equivalent to standard replicator dynamics. However, for our simulations we use a mutation probability $\mu = 0.03$ so

$$Q = (1 - \frac{4}{3}\mu)I + \frac{\mu}{3} \cdot \mathbf{1}. \quad (8)$$

In a large population, the aggregate effective contact rate is determined by the average outcome over all pairs, so $\beta = \frac{\beta_b S_U I_U}{S I}$ where β_b is the baseline effective contact rate, $S_U = S \times (f(UP, t) + f(UU, t))$ is the proportion of susceptible individuals playing U , and $I_U = I \times (f(PU, t) + f(UU, t))$ is the proportion of infected individuals playing U .

The combined model equations with evolutionary behavioral dynamics can be written as

$$\begin{aligned} \dot{S} &= \gamma I - \beta_b S_U I_U, \\ \dot{I} &= \beta_b S_U I_U - \gamma I, \\ \dot{\mathbf{f}} &= s[D_{\mathbf{f}}Q(M\mathbf{f}) - (\mathbf{f}^T M\mathbf{f})\mathbf{f}]. \end{aligned} \quad (9)$$

The system above has six compartments. Of these, only four are strictly necessary as $S + I = 1$ and $\sum_j f(\sigma_j, t) = 1$. In a completely susceptible population, action pairs where both susceptibles play P or U are Bayes-Nash equilibria. In order to restrict the domain of possible initial conditions, we take the risk dominant equilibria, or mixtures between UP and UU . This can be interpreted as individuals in a disease free state preferring unprotected sex because it best hedges against uncertainty in partner actions. While the $\sigma(\theta_i = I)$ actions are never realized at the DFE, the type-contingent strategy framework implies that individuals must be able to specify an action that they would take if they became infected. The effective contact rate for a single infected individual is $(f(UU, 0) + f(PU, 0))\beta_b$,


where $f(UU, 0), f(PU, 0) \in [0, 1]$. As a result, the basic reproductive rate of the combined model is

$$\mathcal{R}_0 = \frac{(f(UU, 0) + f(PU, 0))\beta_b}{\gamma}. \quad (10)$$


Thus, \mathcal{R}_0 depends on the behavioral initial condition, unlike the standard SIS model. The value of the above expression will lie within the interval $[0, \beta_b/\gamma]$. The payoff and risk dominance of the (U, U) Nash equilibrium at the initial disease-free steady state (discussed above) suggests that $f(PU, 0)$ is likely to be low, and indeed for the remainder of this paper we will assume that the initial infected plays the strategy UU , while the underlying susceptible population at $t = 0$ will be assumed to take on a mixture of the UU and UP strategies. It is nonetheless of interest to note that the dependence of \mathcal{R}_0 on the initial behavioral conditions suggests the potential for both high \mathcal{R}_0 with disease extinction and low \mathcal{R}_0 with endemic prevalence: if $f(PU, 0) \approx 1$ it may be possible to have high \mathcal{R}_0 but still have disease extinction due to all susceptibles playing the protected strategy, while conversely if \mathcal{R}_0 is low due to low frequency of UU and PU but otherwise has a high contact rate, it may be possible to generate endemic prevalence even for $\mathcal{R}_0 < 1$ (as explored further below).

3. Results


Due to the addition of the replicator-mutator equations, it is not straightforward to solve for the steady states of the combined model analytically. Instead, we conducted a range of numerical simulations to investigate the effects of behavioral dynamics using Python 2.7 with Numpy 1.9, Scipy 0.15.0, and Matplotlib 1.4.2. In order to explore a range of outcomes generated by the combined model, we focused on both long-term and short-term dynamics. Given the addition of behavior change, the short term dynamics provided useful mechanistic insights to explain the steady state properties of the model. Unless noted otherwise we used $a = 1, b = 0.6, c = 0.4, d = 0.2$ for the values of the outcome payoffs. In general the qualitative features of the model do not change substantially for different payoff values provided the overall scale and ordering are preserved. Unless otherwise specified, the disease initial conditions were $S(0) = 0.99, I(0) = 0.01$. In the first section below, we consider fixed initial behavioral conditions at $f(UU, 0) = f(UP, 0) = 0.5$, while in the following section we consider the effects of varying the initial behavioral conditions (for $f(UU, 0)$ between 0 and 1).




(a)



(b)



(c)



(d)

Figure 2: Simulated trajectories for the combined model with $R_0 = 1.2$. Disease parameters $\beta_b = 1.2, \gamma = 0.5$ with behavioral parameters $a = 1, b = 0.6, c = 0.4, d = 0.2$. (a) Disease prevalence, (b) Type-contingent strategies, (c) Fractional contact reduction, (d) Strategy fitness.

3.1. Model Dynamics

We began by inspecting an example set of model trajectories, shown in Fig. 2 (with $\beta_b = 1.2$, $\gamma = 0.5$, and $f(UU, 0) = 0.5$ to give $\mathcal{R}_0 = 1.2$), to illustrate the interactions between behavioral changes and disease dynamics. Comparing the prevalence and contact reduction trajectories demonstrated the impact of the adaptive behavioral dynamics on the progression of the simulated outbreak. In this example the outbreak was relatively small, so susceptible individuals did not have a large incentive to use protection. However, the initial increase in disease prevalence favored strategies where infected individuals used protection, which in turn decreased the force of infection enough to halt the initial outbreak. However, as prevalence decreased, unprotected strategies became less costly and contact reduction decreased again. Consequently there was a small secondary outbreak. The timing of the behavioral response also appeared to be influential. While contact reduction tracked prevalence over time, it did so at a delay. Thus the level of contact reduction was still relatively high at the onset of the second outbreak, preventing a large secondary peak. This preliminary exploration suggested that both the infectiousness of the disease and the speed of behavioral adaptation play important roles in the overall model dynamics.

To explore this issue further, we defined a timescale parameter s for the behavioral dynamics, given as a scaling factor on the payoff matrix A in Eq. (7). From a behavioral perspective, this parameter can be thought of as controlling the speed of adaptation in the population. We then evaluated the long-term dynamics of the model in steady state for a range of values of \mathcal{R}_0 and s , shown in Fig. 3 (with all remaining parameters as in Fig. 2). In the triangular region about $\mathcal{R}_0 = 2.7$ and the canyon between $3 < \mathcal{R}_0 < 3.5$ the damped oscillations became sustained with a substantially higher average steady state prevalence. The triangular region of oscillations gives way to another endemic steady state region as \mathcal{R}_0 increases, followed by the thin band of oscillations and eventually disease extinction in the triangular region along the right edge.

Figure 4 and Supplementary Figure B.8 provide a more detailed examination of the dynamics in Fig. 3 as \mathcal{R}_0 increases for $s = 1$. For \mathcal{R}_0 sufficiently above 2, the early increase in contact reduction by infected individuals was not sufficient to stop the spread of the initial outbreak (4b). As a result prevalence increased, causing infecteds to begin to switch back to unprotected strategies (B.8 b).




Figure 3: Long term dynamics of the combined model with increasing values of R_0 and the behavioral scale parameter s . R_0 was adjusted by varying the parameter β_b in Eq. (10). (a) The amplitude of steady state prevalence oscillations. (b) The average prevalence at steady state.

However, at higher prevalence susceptible individuals had an incentive to use protection, compensating for the behavior of the infected population. As in the initial example, this overall increase in contact reduction reduced the force of infection and halted the growth of the outbreak. The magnitude of this response appeared to determine the long term dynamics of the model. When prevalence was reduced to a moderate level neither protective nor unprotected strategies were able to gain a stable foothold in the population, leading to sustained oscillations in both contact reduction and prevalence (4b, B.8 b). A stronger contact reduction response pushed the disease to a level where infected individuals again had a strong incentive to use protection. The fitness advantage of this strategy was sufficient to survive the second outbreak, leading to an endemic steady state (4c, B.8 c). Notably, when R_0 was high enough that the resulting behavioral response nearly drove the disease to extinction, sustained oscillations were again possible (4d, B.8 d). Subsequent outbreaks in this scenario were contained by increasing contact reduction from a higher baseline due to the initial outbreak, but grew rapidly due to the infectiousness of the disease. Finally, when the initial outbreak was extremely large, essentially all remaining susceptibles (including those who recently recovered) played protective strategies, driving the disease to extinction (4e, B.8 e). Since the boundaries of both oscillatory regions shifted with both s and R_0 , it appeared that the timescales of infection and adaptation must align in order to produce sustained oscillations (e.g. Figure 4b and B.8b). In particular, oscillatory dynamics appear only to be possible for relatively

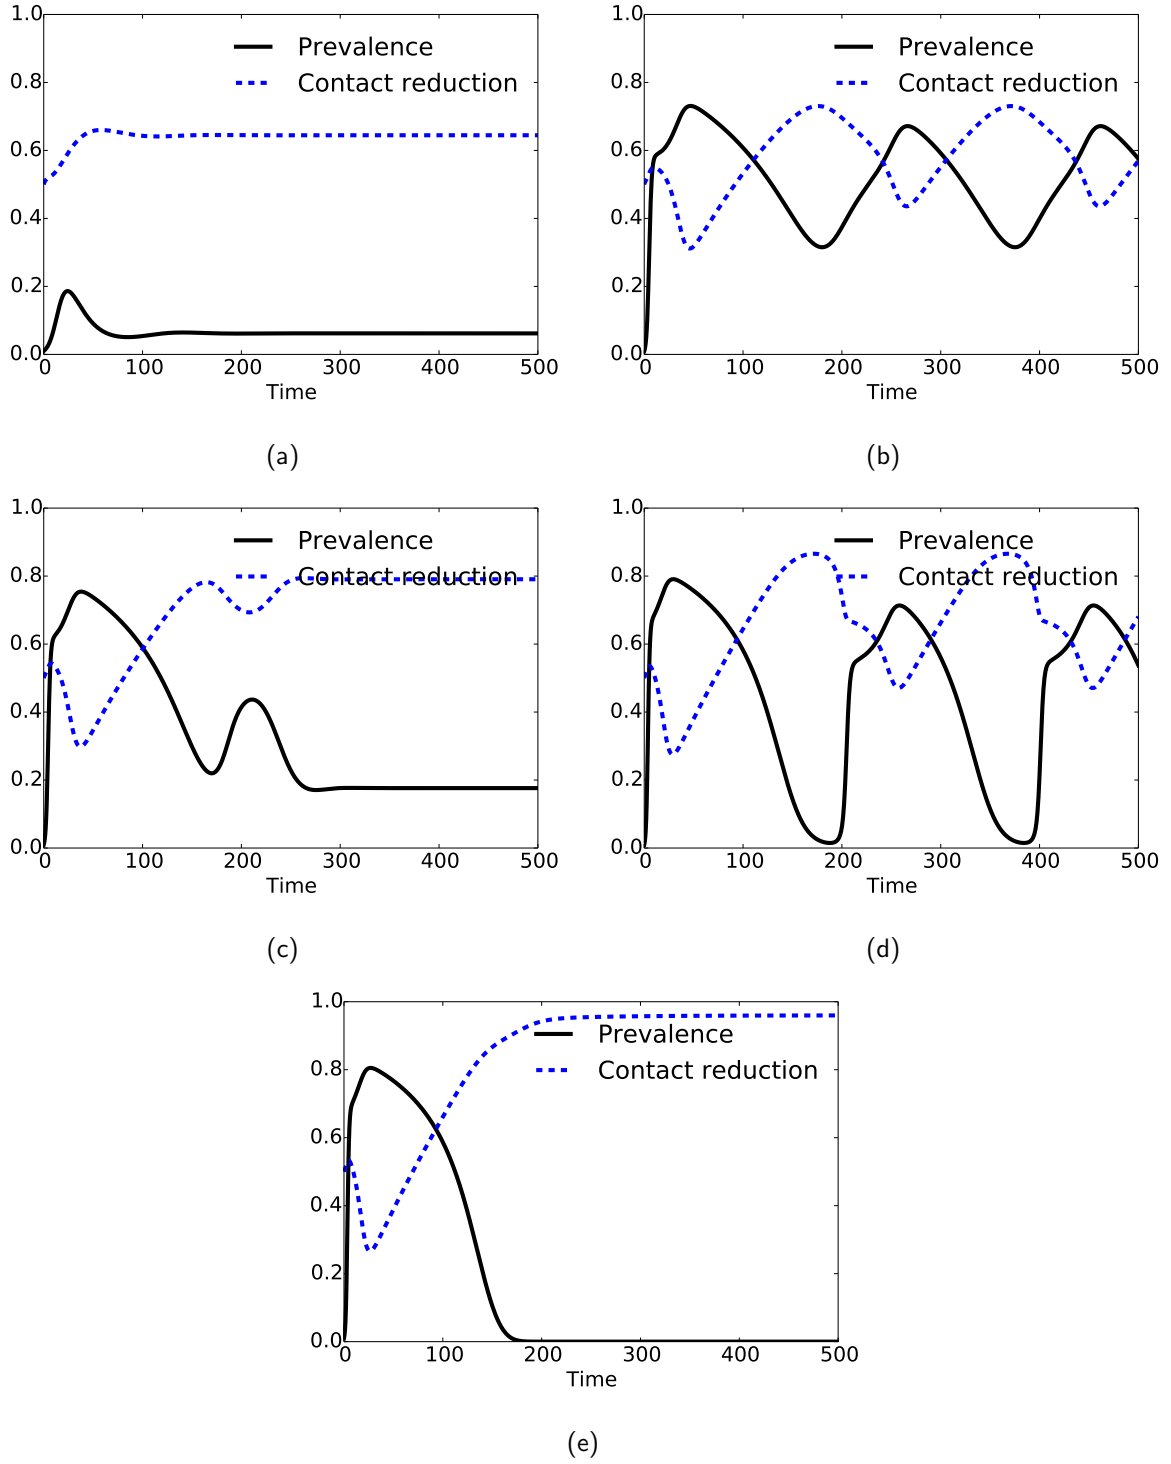


Figure 4: Prevalence and contact reduction from the combined model with $\gamma = 0.5$, $\mathbf{f}(0) = (0.0, 0.0, 0.5, 0.5)$ for all simulations and (a) $\beta_b = \mathcal{R}_0 = 1.5$, (b) $\mathcal{R}_0 = 2.7$, (c) $\mathcal{R}_0 = 2.9$, (d) $\mathcal{R}_0 = 3.3$, and (e) $\mathcal{R}_0 = 3.5$.




Figure 5: Average steady state prevalence of the combined model for increasing \mathcal{R}_0 with $\gamma = 0.5, \mathbf{f}(0) = (0.0, 0.0, 0.5, 0.5)$. The transition between the disease free and endemic equilibrium occurs at $1 < \mathcal{R}_0 < 1.05$.

slow adaptation speeds, suggesting that the lag induced in the behavioral dynamics contributes to the potential for oscillations.

Using the same parameters as in Fig. 4, in Figure 5 we also examined steady state prevalence near $\mathcal{R}_0 = 1$, to evaluate whether behavior changes may affect the threshold for generating an outbreak. We found that steady state prevalence remained zero even when \mathcal{R}_0 was greater than 1, suggesting that behavioral dynamics were able to extinguish the disease even when the initial growth rate would have generated and epidemic.

3.2. Behavioral Initial Condition Analysis

For our next set of simulations, we tested the effect of varying the initial behavioral conditions. Figure 6a shows the average steady state prevalence across $f(UU, 0) \in [0, 1]$ when $\beta_b/\gamma = 3$ (where the remaining population was assumed to take strategy UP as noted above). As suggested by the form of \mathcal{R}_0 , an endemic steady state persisted even for $\mathcal{R}_0 < 1$. This counterintuitive phenomenon occurred because, while the disease initially declined, it did not immediately become extinct (Fig. C.9). The declining prevalence led to increased adoption of UU , driving the effective reproductive rate back above one. Figure 6b depicts the effect of different levels of $f(UU, 0)$ on the steady state behavior of the combined model for increasing β_b/γ . A wide range of dynamics were observed as β_b/γ varied, including a single stable equilibrium, bistability, and oscillations. When $f(UU, 0)$ was low (green line in Fig. 6b), the model only exhibited damped oscillations, and did not produce the same extinction behavior at high




Figure 6: The effect of changing the initial condition $f(UU,0)$ with $\gamma = 0.5$. (a) Steady state prevalence and \mathcal{R}_0 for $\beta_b/\gamma = 3$ (Dashed line indicates where $\mathcal{R}_0 = 1$). (b) Steady state prevalence of the combined model as β_b/γ increases for three choices of $f(UU,0)$. In regions of β_b/γ where sustained oscillations occurred, the solid black line shows the value of the crest while the dashed line shows the value of the trough.

β_b/γ as in the case where $f(UU,0)$ was substantially larger than zero. For $f(UU,0) = 0.5$ (blue dotted line), the model exhibited the same bifurcation pattern as in the previous section. At high $f(UU,0)$ (black line) the behavior was similar, however the endemic prevalence before the first oscillatory region was substantially greater. From this, $f(UU,0)$ appeared to act as a switch between possible steady state regimes of long term dynamics.

For the constant steady state regions, the endemic prevalence took one of two values for a given β_b/γ depending on $f(UU,0)$, similar to Fig. 6a. More broadly, for any fixed β_b/γ , there were two basins of attraction corresponding to either the $f(UU,0) = 0$ or $f(UU,0) = 1$ case, where the unstable equilibrium dividing the two regions depended on the value of β_b/γ . This is illustrated by the $f(UU,0) = 0.5$ curve in Fig. 6b, which switches between existing in the $f(UU,0) = 0$ and $f(UU,0) = 1$ basins around $\beta_b/\gamma = 5$. There are also regions of Fig. 6b where only a single steady state exists regardless of the value of $f(UU,0)$, at very low values of β_b/γ (left portion just above $\beta_b/\gamma = 2$) and larger values (between the oscillatory regions).

As noted in the previous section, there were also multiple regions where the long term dynamics showed disease extinction even though $\mathcal{R}_0 > 1$. For $f(UU,0) = 0.5$ and 1, the left corner of Fig. 6b shows extinction with $\mathcal{R}_0 > 1$, and the right corner shows the same extinction discussed in the previous

section for the $f(UU, 0) = 1$ basin at high values of β_b/γ .

3.3. Comparison with Alternate Models

Finally, we considered how the dynamics of the model compare to other potential models of the disease dynamics, to examine how neglecting the behavioral dynamics may alter model forecasts of the epidemic trajectory or affect estimates of key epidemiological parameters such as \mathcal{R}_0 . As a first example, we chose a fixed contact rate SIS model parameterized such that the \mathcal{R}_0 of both models was identical. While the trajectories were similar at very early times (at which the epidemic growth rate can still be characterized by \mathcal{R}_0), the combined model quickly diverged due to the reduced effective contact rate, and equilibrated at a substantially lower endemic level (Figure 7a). Similarly, we computed $\mathcal{R}_0 = \frac{1}{1-I_\infty}$ naively from the steady state prevalence of the combined model in Figure 7a, without accounting for a time-varying contact rate. The estimated value of 1.03 was substantially lower than the true \mathcal{R}_0 of 1.2 for the combined model. This difference can be even more severe when considering endemic steady states at higher \mathcal{R}_0 's (such as the area to the right of the triangular region in Figure 3). For example, the endemic steady state in Figure 4c yielded an apparent \mathcal{R}_0 of 1.21, when the underlying \mathcal{R}_0 for the combined model was 2.9. Given the simplicity of the SIS model, these discrepancies were not surprising. Nonetheless, they highlight some pitfalls of neglecting the effects of behavioral dynamics.

However, as the SIS model is known not to be able to produce oscillations, it might not be a likely choice given data that came from real-world infection dynamics similar to the combined model. Due to the delay between recovering from infection and returning to susceptibility, the SIRS model is capable of producing damped oscillations, and so might make a reasonable initial guess for the model structure if the length of immunity was unknown. To evaluate how a more realistic but misspecified model might affect parameter estimation, we fit the SIRS model (Appendix D) to the prevalence trajectory of the combined model with $\mathcal{R}_0 = 1.2$, using least squares with Nelder-Mead optimization, as shown in Figure 7b. While the best fit SIRS model conformed well to the target trajectory, the resulting parameter estimates included both a four fold decrease in the baseline contact rate and a substantial increase in the average waiting time until a recovered individual becomes susceptible again. The \mathcal{R}_0 for the best fit SIRS model was 1.35, a 12.3% increase compared to the behavioral model. In this case, the feedback between behavior and transmission may make interventions more effective than would appear to be the

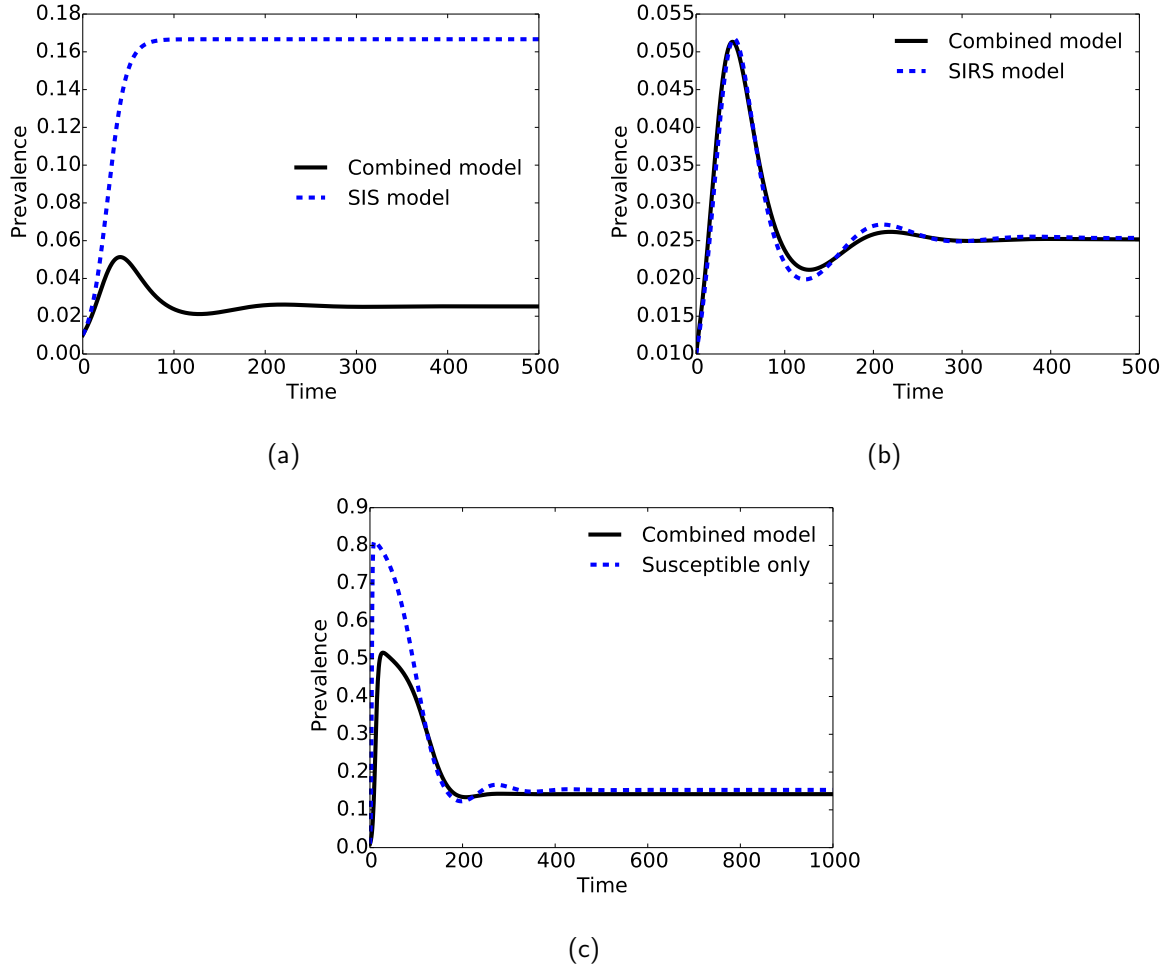


Figure 7: Comparison of simulations of the combined model other compartmental models. (a) SIS model contact rate $\beta = 0.6, \gamma = 0.5$ to match the $R_0 = 1.2$ of the combined model, (b) Prevalence trajectories for the combined model and an SIRS model fit to the combined model prevalence simulated with $R_0 = 1.2$. SIRS best fit parameters: $\beta = 0.274, \gamma = 0.203, \delta = 0.022$. (c) A variant combined model in which only susceptible individuals adapt their behavior fit to the simulated trajectory from the full model with $\beta_b = 1.8, \gamma = 0.5$ and $R_0 = 2.7$. The best fit parameter $\beta_b^* = 2.84$ for the reduced model, giving an $R_0 = 5.68$.

case if we had assumed a fixed contact rate.

Finally, most previous work combining game theory with transmission models to model contact reduction has focused on the adaptive behavior of susceptibles [14, 18, 23, 24]. This is a natural formulation for vaccination, but may be less suitable for more general contact reduction behaviors. Thus, to test the effect of modeling adaptive behavior by both susceptible and infected individuals, we used a reduced variant of our combined model in which infected individuals always select the action U . As with the SIRS model, we fit this reduced model to a simulated trajectory from the full model. In this case, the reduced model overshoots the initial outbreak curve, but equilibrates to a similar endemic equilibrium to the full model (Figure 7c). However, the best fit parameter and \mathcal{R}_0 for the reduced model were nearly two times higher than the true values for the full model. For a wider range of parameter values of the full model, the reduced model generally failed to provide qualitatively accurate fits. In addition, the reduced model did not produce the same range of dynamics as the full model. We give a full description of the reduced model in Appendix D as well as an expanded set of fits and simulations.

4. Discussion

To explore the feedback between behavior and disease dynamics, we developed a model of disease transmission with adaptive behavior among both susceptible and infected individuals. Numerical simulations illustrated the effect of behavior-disease feedback on model dynamics and inferences about STI transmission (Figure 2). We found that phase transitions between damped and sustained oscillations occurred at an intermediate transmission rate and again at a higher transmission rate (Figure 3). Increasing the adaptation rate parameter s reduced the range of transmission rates at which sustained oscillations could occur.

Disease extinction occurred both for low and high transmission rates (Figures 3, 6). At lower transmission rates, adoption of protection by infected individuals lead to extinction above the typical $\mathcal{R}_0 = 1$ threshold (Figure 5). This suggests that behavior-disease feedback can create a herd-protection-like effect if the infected population can be reduced sufficiently to incentivize the use of protection among that group. At high transmission rates and low-to-mid range levels of initial risky behavior ($f(UU, 0)$) in the infected population, we observed the counterintuitive result that a large initial increase in risky

behavior (UU and PU strategies) among infected individuals preceded disease extinction. However, because susceptible individuals were also switching to protective strategies (PP and PU) rapidly, the combined effect of both susceptibles and infected adopting PU resulted in sufficient growth of PU to drive the disease to extinction as the effective number of susceptible individuals remained small and all were largely playing protective strategies (Figure 4, B.8). In fact, both the high \mathcal{R}_0 and low \mathcal{R}_0 extinction phenomena (for $\mathcal{R}_0 > 1$ but still low) appear to be driven by a reduction in the effective size of the ‘bottleneck’ population group. The low \mathcal{R}_0 extinction (Figure 5) seems largely due to increasing infected protective strategies, while the high \mathcal{R}_0 extinction seen on the right side of Figure 4 is largely due to increasing susceptible protective strategies. In each case, the low/high prevalence makes infecteds/susceptibles a bottleneck group, where sufficient protective response in that group appears to result in disease elimination. The feedback between disease and behavior thus often constrains endemic outcomes to scenarios where a disease is only moderately infectious, for a range of realistic levels of initial risky behavior (Figure 6b).

However, simulations predicted a more pessimistic outcome for changes in the initial strategic distribution. In general, the initial level at which infected individuals used protection did not determine whether the disease would become extinct. As long as the value of β_b/γ was sufficiently large, any $f(UU, 0)$ led to an outbreak and endemic disease (Figure 6). Instead, the initial distribution acted as a switch between two steady state regimes: one in which initial risky behavior by infecteds (high $f(UU, 0)$) lead to sustained oscillations and higher endemic prevalence could occur for values of c_b below the first oscillatory region, and another in which initial protective behavior by infecteds (low $f(UU, 0)$) yielded damped oscillations and generally lower endemic prevalence. The exception to this pattern is the region where β_b/γ is large. As noted above, in this region extinction occurs when early risky behavior by infected individuals is then countered by widespread use of protection by susceptible individuals. However, early widespread protection use by infecteds leads to a smaller outbreak that is slowed, but not eliminated by susceptible protection use. This complicates the interpretation of \mathcal{R}_0 for the combined model, as the standard endemic threshold property often does not apply. For interventions it is thus important to understand the interaction between early behavioral patterns, the transmission rate, and treatment levels. Interventions based on these factors may yield counterintuitive

results depending on the other factors, making comprehensive data collection critical when optimizing intervention strategies.

It is also interesting to note that while an analysis of the Bayes-Nash equilibria of the protected sex game at DFE indicates that the susceptible-unprotected strategies (U, U) are risk dominant, our combined model suggests that the actual initial distribution will depend on a population's previous experience with disease. That is, when the outbreak is large (i.e. for high \mathcal{R}_0 , Figure 4), adoption of protective strategies results in disease elimination. However, the increased protective strategies by susceptibles mean that the model switches to the attraction basin of the non-dominant (P, P) Nash equilibrium as it goes to the new disease-free steady state B.8 (or as close as possible given the mutation rates in the replicator-mutator equation). By contrast, below the threshold for high \mathcal{R}_0 extinction, the model tends to reach a steady state close to the susceptible-unprotected (U, U) Nash equilibrium since endemic prevalence is low enough to favor susceptible-susceptible contact. This suggests that in general, the disease-free strategy balance between protected and unprotected depends on the magnitude of previous outbreaks, with larger outbreaks yielding disease-free protective behaviors even though they are neither payoff nor risk dominant. As a result, post-outbreak surveillance and control are crucial to prevent recurrent outbreaks due to reintroduction.

Based on comparisons between the combined model and similar disease models with simplified contact dynamics, predictions from fixed-contact rate models may omit important dynamical features or yield misleading parameter estimates (Figure 7). Similarly, if an STI model is misspecified, the common method of computing \mathcal{R}_0 from endemic prevalence when a disease is assumed to be at equilibrium [41] is not appropriate and yields inaccurate estimates. One advantage of using a game theoretic framework to model interactive behavior is the flexibility of ordinal utility. Since protective measures can be used by susceptible or infected individuals to prevent infection or transmission, respectively, we used a preference order that gave individuals in both disease states an incentive to use protection if paired with a partner of the opposite state. Our simulation results suggest that disease dynamics differ in this scenario as opposed to when only susceptible individuals adapt their behavior, and indeed model fits using the reduced model considering only susceptible behavior change often yielded both incorrect \mathcal{R}_0 estimates and dynamic trajectories (Figure D.11). The role of infected behavior dynamics was

particularly notable in the early stage of outbreaks when disease was not prevalent enough to induce susceptible individuals to use protection. In addition, the shift between infected and susceptible adoption of protective strategies appeared to drive sustained oscillations. To more thoroughly consider the effects of infected behavior dynamics, we also considered an alternative preference structure in which infected individuals no longer explicitly prefer protected sex with susceptible individuals, but rather have the same preferences regardless of their partner's disease status (Appendix E). This model only exhibits one oscillatory region and a higher endemic prevalence than the original model (Fig. E.12). However, extinction still occurs at high \mathcal{R}_0 , confirming that susceptible protective strategies drive the high- \mathcal{R}_0 extinction phenomenon.

In order to focus on the effect of adaptive protective behavior, our combined model simplifies many other realistic factors that contribute to STI transmission. The protected sex game does not include explicit negotiation, which likely biases the effective contact rate downward, since negotiation or coercion could lead to unprotected sex even when the initial action pair is (U, P) . The combined model presented here still uses mass action assumptions to determine interactions between individuals. Although more complex contact patterns are known to influence transmission dynamics, we model a homogeneous population to focus on the effect of time-varying behavior. While this may not be completely implausible in a highly active group such as MSM frequenting bathhouses, it is almost certainly a poor representation of the manner in which individuals form sexual partnerships. Typically, adding contact heterogeneity increases the \mathcal{R}_0 of a model, so we might expect more rapid early outbreak growth for a wider range of parameters in our model, potentially expanding the regions in which sustained oscillations or even disease extinction occur. Our general framework, however, is amenable to extensions to expand the state space of the disease model to represent more complex natural histories or population structures as with standard transmission models. In addition, it is possible to model the combined process on a contact network or using a stochastic framework. One particular extension that could yield insight into spatial patterns and the effect of local information would be a simulation of the model on a regular lattice, where individuals could use either a global or local prior to estimate the probability of their partner's type.

A particularly compelling consequence of developments in economic-epidemiological models is the

potential to estimate more complex behavioral parameters using traditional surveillance data. In the context of our combined model, this would result in estimates for the relative payoff values (a, b, c, d) , which capture useful information about the perceptions of populations facing STI outbreaks. While subject to as yet unknown identifiability properties, this manner of parameter estimation could provide a valuable link between game theoretic methods and the extensive empirical literature on the epidemiology of STIs.

Acknowledgements

This work was partially supported by a University of Michigan Block Grant to MALH.

References

References

- [1] National Center for HIV/AIDS, Viral Hepatitis, STD, and TB Prevention, Fact sheet: Incidence, prevalence, and the cost of sexually transmitted infections in the United States, Tech. rep., Centers for Disease Control & Prevention (2013).
- [2] World Health Organization Media Centre, Sexually transmitted infections, Tech. rep., World Health Organization (2013).
- [3] R. M. Anderson, G. P. Garnett, Mathematical models of the transmission and control of sexually transmitted diseases, *Sexually transmitted diseases* 27 (10) (2000) 636–643.
- [4] K. T. Eames, M. J. Keeling, Modeling dynamic and network heterogeneities in the spread of sexually transmitted diseases, *Proceedings of the National Academy of Sciences* 99 (20) (2002) 13330–13335.
- [5] A. G. Lam, A. Mak, P. D. Lindsay, S. T. Russell, What really works? an exploratory study of condom negotiation strategies, *AIDS Education and Prevention* 16 (2) (2004) 160–171.
- [6] C. Peasant, G. R. Parra, T. M. Okwumabua, Condom negotiation: Findings and future directions, *The Journal of Sex Research* (ahead-of-print) (2014) 1–14.
- [7] D. C. Des Jarlais, P. Friedmann, H. Hagan, S. R. Friedman, The protective effect of AIDS-related behavioral change among injection drug users: a cross-national study. WHO multi-centre study of AIDS and injecting drug use., *American Journal of Public Health* 86 (12) (1996) 1780–1785.
- [8] G. K. SteelFisher, R. J. Blendon, M. M. Bekheit, K. Lubell, The public's response to the 2009 H1N1 influenza pandemic, *New England Journal of Medicine* 362 (22).
- [9] G. J. Rubin, R. Amlôt, L. Page, S. Wessely, et al., Public perceptions, anxiety, and behaviour change in relation to the swine flu outbreak: cross sectional telephone survey, *Bmj* 339.

- [10] J. D. Fisher, W. A. Fisher, Changing aids-risk behavior, *Psychological bulletin* 111 (3) (1992) 455.
- [11] S. Gregson, G. P. Garnett, C. A. Nyamukapa, T. B. Hallett, J. J. Lewis, P. R. Mason, S. K. Chandiwana, R. M. Anderson, HIV decline associated with behavior change in eastern Zimbabwe, *Science* 311 (5761) (2006) 664–666.
- [12] E. P. Fenichel, C. Castillo-Chavez, M. Ceddia, G. Chowell, P. A. G. Parra, G. J. Hickling, G. Holloway, R. Horan, B. Morin, C. Perrings, et al., Adaptive human behavior in epidemiological models, *Proceedings of the National Academy of Sciences* 108 (15) (2011) 6306–6311.
- [13] K. D. Schroeder, F. G. Rojas, A game theoretical analysis of sexually transmitted disease epidemics, *Rationality and Society* 14 (3) (2002) 353–383.
- [14] C. T. Bauch, Imitation dynamics predict vaccinating behaviour, *Proceedings of the Royal Society B: Biological Sciences* 272 (1573) (2005) 1669–1675.
- [15] C. T. Bauch, D. J. Earn, Vaccination and the theory of games, *Proceedings of the National Academy of Sciences of the United States of America* 101 (36) (2004) 13391–13394.
- [16] T. C. Reluga, A. P. Galvani, A general approach for population games with application to vaccination, *Mathematical biosciences* 230 (2) (2011) 67–78.
- [17] T. C. Reluga, C. T. Bauch, A. P. Galvani, Evolving public perceptions and stability in vaccine uptake, *Mathematical biosciences* 204 (2) (2006) 185–198.
- [18] T. C. Reluga, Game theory of social distancing in response to an epidemic, *PLoS computational biology* 6 (5) (2010) e1000793.
- [19] D. Fudenberg, J. Tirole, *Game theory*, MIT Press, 1991.
- [20] R. Selten, Bounded rationality, *Journal of Institutional and Theoretical Economics (JITE)/Zeitschrift für die gesamte Staatswissenschaft* (1990) 649–658.
- [21] D. Kahneman, New challenges to the rationality assumption, *Journal of Institutional and Theoretical Economics (JITE)/Zeitschrift für die gesamte Staatswissenschaft* (1994) 18–36.
- [22] M. H. Bazerman, Negotiator judgment a critical look at the rationality assumption, *American Behavioral Scientist* 27 (2) (1983) 211–228.
- [23] F. H. Chen, Rational behavioral response and the transmission of stds, *Theoretical population biology* 66 (4) (2004) 307–316.
- [24] P.-Y. Geoffard, T. Philipson, Rational epidemics and their public control, *International Economic Review* (1996) 603–624.
- [25] R. J. Wolitski, C. J. Bailey, A. O'Leary, C. A. Gómez, J. T. Parsons, Self-perceived responsibility of HIV-seropositive men who have sex with men for preventing hiv transmission, *AIDS and Behavior* 7 (4) (2003) 363–372.
- [26] A. O'Leary, Guessing games: Sex partner serostatus assumptions among hiv-positive gay and bisexual men.
- [27] C. E. Kennedy, A. M. Medley, M. D. Sweat, K. R. O'Reilly, Behavioural interventions for HIV positive prevention in developing countries: a systematic review and meta-analysis, *Bulletin of the World Health Organization* 88 (8) (2010) 615–623.

- [28] C. J. De Rosa, G. Marks, Preventive counseling of HIV-positive men and self-disclosure of serostatus to sex partners: New opportunities for prevention., *Health Psychology* 17 (3) (1998) 224.
- [29] R. Breban, R. Vardavas, S. Blower, Mean-field analysis of an inductive reasoning game: application to influenza vaccination, *Physical review E* 76 (3) (2007) 031127.
- [30] J. M. Smith, *Evolution and the Theory of Games*, Cambridge university press, 1982.
- [31] J. Hofbauer, K. Sigmund, Evolutionary game dynamics, *Bulletin of the American Mathematical Society* 40 (4) (2003) 479–519.
- [32] R. Anderson, Transmission dynamics and control of infectious disease agents, in: *Population biology of infectious diseases*, Springer, 1982, pp. 149–176.
- [33] R. M. Anderson, R. M. May, *Infectious diseases of humans*, Vol. 1, Oxford university press Oxford, 1991.
- [34] J. A. Yorke, H. W. Hethcote, A. Nold, Dynamics and control of the transmission of gonorrhea, *Sexually transmitted diseases* 5 (2) (1978) 51–56.
- [35] J. C. Harsanyi, Games with incomplete information played by bayesian players, i-iii part i. the basic model, *Management science* 14 (3) (1967) 159–182.
- [36] B. J. Kaplan, V. T. Shayne, Unsafe sex: decision-making biases and heuristics., *AIDS Education and Prevention*.
- [37] R. S. Gold, M. J. Skinner, Judging a book by its cover: gay men's use of perceptible characteristics to infer antibody status, *International journal of STD & AIDS* 7 (1) (1996) 39–43.
- [38] S. Z. Wiktor, R. J. Biggar, M. Melbye, P. Ebbesen, G. Colclough, R. DiGioia, W. C. Sanchez, R. J. Grossman, J. J. Goedert, Effect of knowledge of human immunodeficiency virus infection status on sexual activity among homosexual men., *JAIDS Journal of Acquired Immune Deficiency Syndromes* 3 (1) (1990) 62–68.
- [39] J. M. Dawson, R. M. Fitzpatrick, G. Reeves, M. Boulton, J. McLean, G. J. Hart, M. Brookes, et al., Awareness of sexual partners' HIV status as an influence upon high-risk sexual behaviour., *Aids* 8 (6) (1994) 837–842.
- [40] K. M. Page, M. A. Nowak, Unifying evolutionary dynamics, *Journal of theoretical biology* 219 (1) (2002) 93–98.
- [41] J. Heffernan, R. Smith, L. Wahl, Perspectives on the basic reproductive ratio, *Journal of the Royal Society Interface* 2 (4) (2005) 281–293.
- [42] S. Tadelis, *Game theory an introduction*, Princeton university press, 2013.
- [43] N. S. Wenger, F. S. Kusseling, K. Beck, M. F. Shapiro, Sexual behavior of individuals infected with the human immunodeficiency virus: the need for intervention, *Archives of Internal Medicine* 154 (16) (1994) 1849–1854.

Appendix A. Bayesian Games

This section provides a brief overview of static Bayesian games of incomplete information. Those desiring a more complete treatment may refer to [19],[42],and [35].

Appendix A.1. Definition

A normal form n -player symmetric static Bayesian game includes

- A set of players $N = \{1, 2, \dots, n\}$.
- Actions $a_i \in A$ for each player.
- Types $\theta_i \in \Theta$ for each player.
- Pure type-contingent strategies $\sigma : \Theta \rightarrow A$. By convention, we use the notation $\sigma_j = \sigma_j(\theta_i^1)\sigma_j(\theta_i^2)\dots\sigma_j(\theta_i^m)$ to represent the j th strategy for a finite type space.
- Belief distributions p_i where $p_i(\theta_{-i}|\theta_i)$ is the conditional distribution on the types of other players given player i 's type. $-i$ denotes the set of players except i .
- Type-dependent payoffs $u_i : A^n \times \Theta^n \rightarrow \mathbb{R}$ for each player. The expected type-dependent payoff ($E[u_i(a_1, \dots, a_n|\theta_i)]$) gives the average payoff over player i 's belief regarding the other players' types conditional on player i 's own type. $E(u_i(a_1, \dots, a_n, \theta_1, \dots, \theta_n))$ is player i 's unconditional average payoff over all n player types, sometimes denoted $E(u_i(\sigma_{j_1}, \dots, \sigma_{j_n}))$ in terms of strategies.

In general, a *profile* of types, actions, or strategies is defined to be a listing of the particular types/actions/strategies (respectively) assigned to each player. While the actions in a static Bayesian game take place simultaneously, it is useful to break the game down into stages as follows

1. Nature chooses a profile of types $(\theta_1, \theta_2, \dots, \theta_n)$ from the common prior distribution.
2. Each player picks a strategy σ_{j_i}
3. Each player learns only his type θ_i .
4. Using Bayes' theorem and the common prior, each player forms beliefs $p_i(\theta_{-i}|\theta_i)$ over other players' types.

5. Players choose actions simultaneously according to their strategy ($a_1 = \sigma_{j_1}(\theta_1)$) to form a profile (a_1, a_2, \dots, a_n) .
6. Players receive their payoffs $u_i(a_1, \dots, a_n, \theta_1, \dots, \theta_n)$ based on the action profile and the type profile.

Appendix A.2. Bayesian Nash Equilibrium

The strategy profile $(\sigma_{j_1}^*, \dots, \sigma_{j_n}^*)$ is a Bayesian Nash equilibrium if for all players and all types, $\sigma_{j_i}^*$ satisfies

$$\sum_{\theta_{-i}} p_i(\theta_{-i}|\theta_i) u_i(\sigma_{j_i}^*(\theta_i), \sigma_{j_{-i}}^*(\theta_{-i}), \theta_i, \theta_{-i}) \geq \sum_{\theta_{-i}} p_i(\theta_{-i}|\theta_i) u_i(\sigma_{k_i}(\theta_i), \sigma_{j_{-i}}^*(\theta_{-i}), \theta_i, \theta_{-i}) \quad (\text{A.1})$$

for any $\sigma_{k_i} \neq \sigma_{j_i}^*$. Equivalently

$$\sum_{\theta_i} Pr(\theta_i) \left[\sum_{\theta_{-i}} p_i(\theta_{-i}|\theta_i) u_i(\sigma_{j_i}^*(\theta_i), \sigma_{j_{-i}}^*(\theta_{-i}), \theta_i, \theta_{-i}) \right] \geq \sum_{\theta_i} Pr(\theta_i) \left[\sum_{\theta_{-i}} p_i(\theta_{-i}|\theta_i) u_i(\sigma_{k_i}(\theta_i), \sigma_{j_{-i}}^*(\theta_{-i}), \theta_i, \theta_{-i}) \right] \quad (\text{A.2})$$

Appendix B. Behavioral Trajectories for Simulations in Figure 4

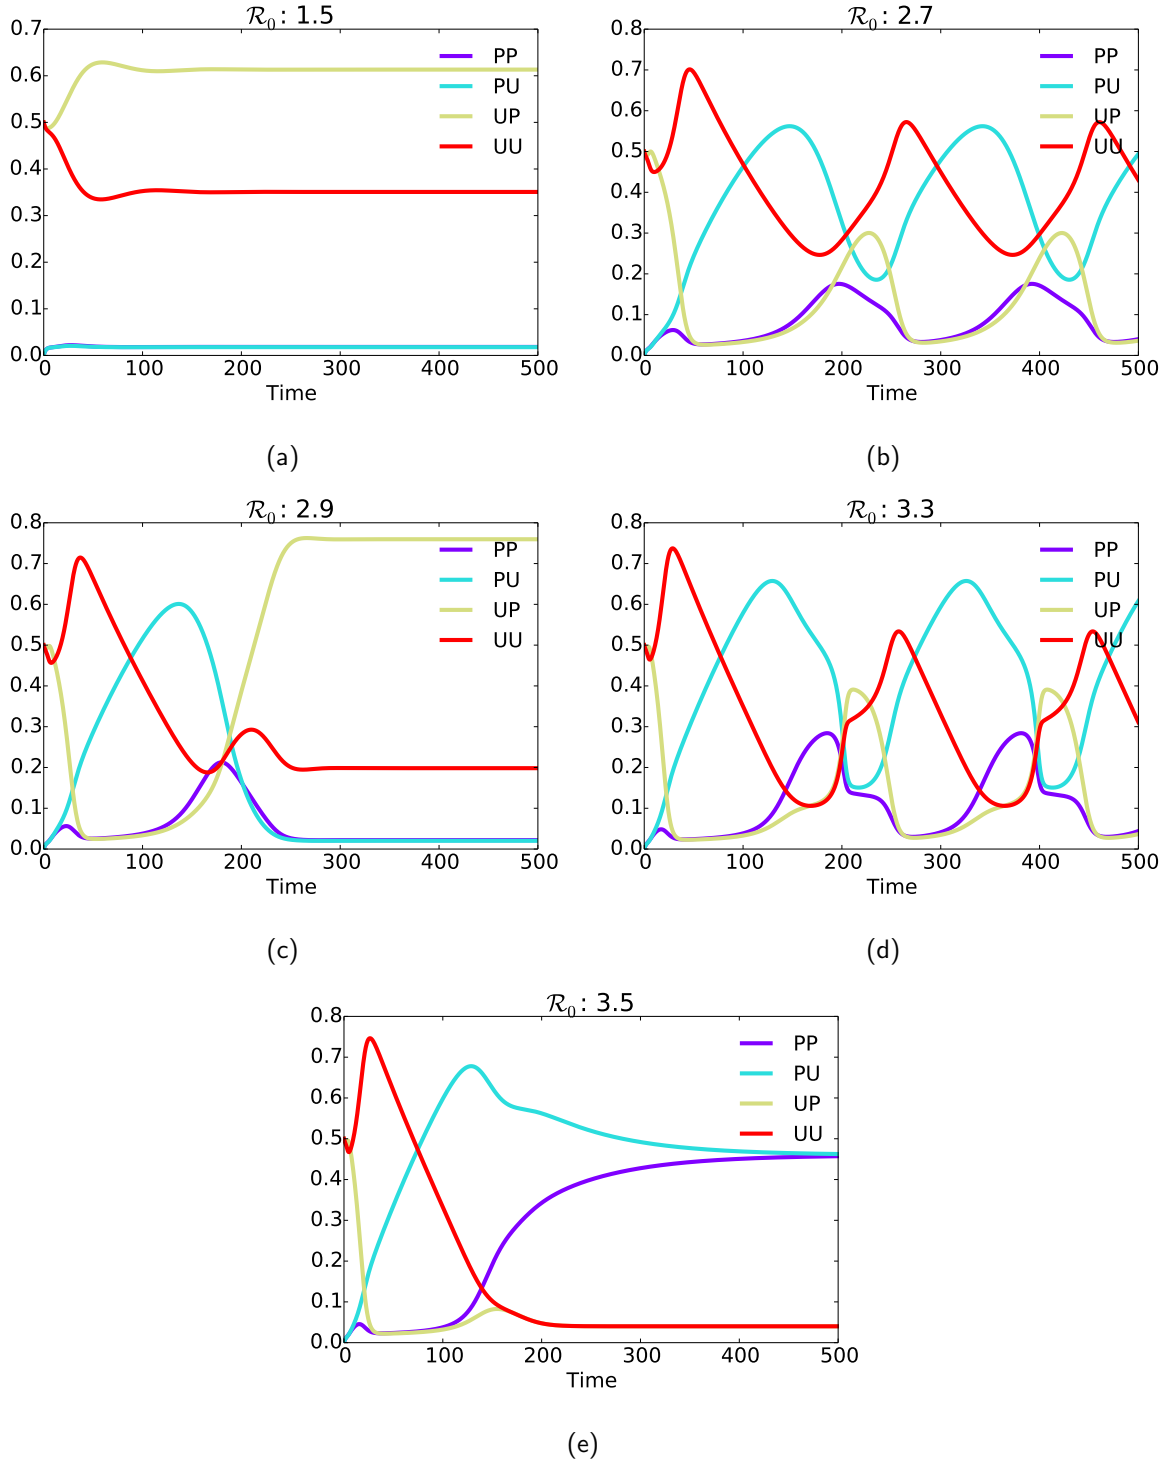


Figure B.8: The distribution of type-contingent strategies as a function of time with $f(UU, 0) = 0.5, \gamma = 0.5$ as in Figure 4.

Appendix C. Endemic Prevalence at $\mathcal{R}_0 < 1$




Figure C.9: Prevalence and contact reduction from the combined model for $\beta_b = 2.1, \gamma = 0.5, f(UU, 0) = 0$.

Appendix D. Alternate Models

Appendix D.1. SIRS Model

$$\begin{aligned}
 \dot{S} &= \delta R - \beta SI \\
 \dot{I} &= \beta SI - \gamma I \\
 \dot{R} &= \gamma I - \delta R
 \end{aligned} \tag{D.1}$$

Appendix D.2. Susceptible-only Behavior Change

For the reduced model, we use the fitness functions

$$\begin{aligned}
 \phi(P_S, t) &= Pr(S)[f(P_S, t)u_1(P_S, P_S, S, S) + f(U_S, t)u_1(P_S, U_S, S, S)] \\
 &\quad + Pr(I)[f(P_I)u_1(P_S, P_I, S, I) + f(U_I)u_1(P_S, U_I, S, I)] \\
 \phi(U_S, t) &= Pr(S)[f(P_S, t)u_1(U_S, P_S, S, S) + f(U_S, t)u_1(U_S, U_S, S, S)] \\
 &\quad + Pr(I)[f(P_I)u_1(U_S, P_I, S, I) + f(U_I)u_1(U_S, U_I, S, I)]
 \end{aligned} \tag{D.2}$$

where the notation P_S denotes a susceptible player choosing P (similarly for U_S, P_I, U_I) and $f(U_I)$ and $F(P_I)$ are fixed over time. Since only susceptible players change their strategy, $\mathbf{f}(t) = (f(P_S, t), f(U_S, t))$, the mutation matrix for this game is 2×2 and we can use the two-dimensional system

$$\begin{aligned}\dot{I} &= -\beta_b S_U f(U_I) I + \gamma I \\ \dot{f}(U_S) &= q_{U_S P_S} \phi(P_S, t) f(P_S, t) + q_{U_S, U_S} \phi(U_S, t) f(U_S, t) - \bar{\phi} f(U_S, t)\end{aligned}\tag{D.3}$$

where $S_U = S f(U_S, t)$. Figure D.10 shows prevalence trajectories for this model at increasing baseline effective contact rates. Figure D.11 shows the best fit trajectories and parameter values for the reduced model compared to the full model.




Figure D.10: Prevalence dynamics for the reduced model from Section 2 where only susceptible individuals adapt their behavior. For the simulations above $\gamma = 0.5$, $\mathbf{f}(0) = (0.0, 1.0)$.

Appendix E. Alternate preferences

We examined a case where infected individuals prefer unprotected sex over protected sex regardless of partner type, a phenomenon that has been observed empirically [43]. For susceptible individuals, the type-dependent payoffs from the alternate game are the same as in Figure 1. Without loss of generality, for an infected-type player 1 paired with a susceptible-type player 2, the type-dependent payoff matrix is

		Player 2	
		P	U
Player 1	P	b	c
	U	c	a

Figure E.12 shows the long-term behavior of the alternate model for a range of \mathcal{R}_0 and s . This model only exhibits one oscillatory region and a higher endemic prevalence than the original model. However, extinction still occurs at high \mathcal{R}_0 .

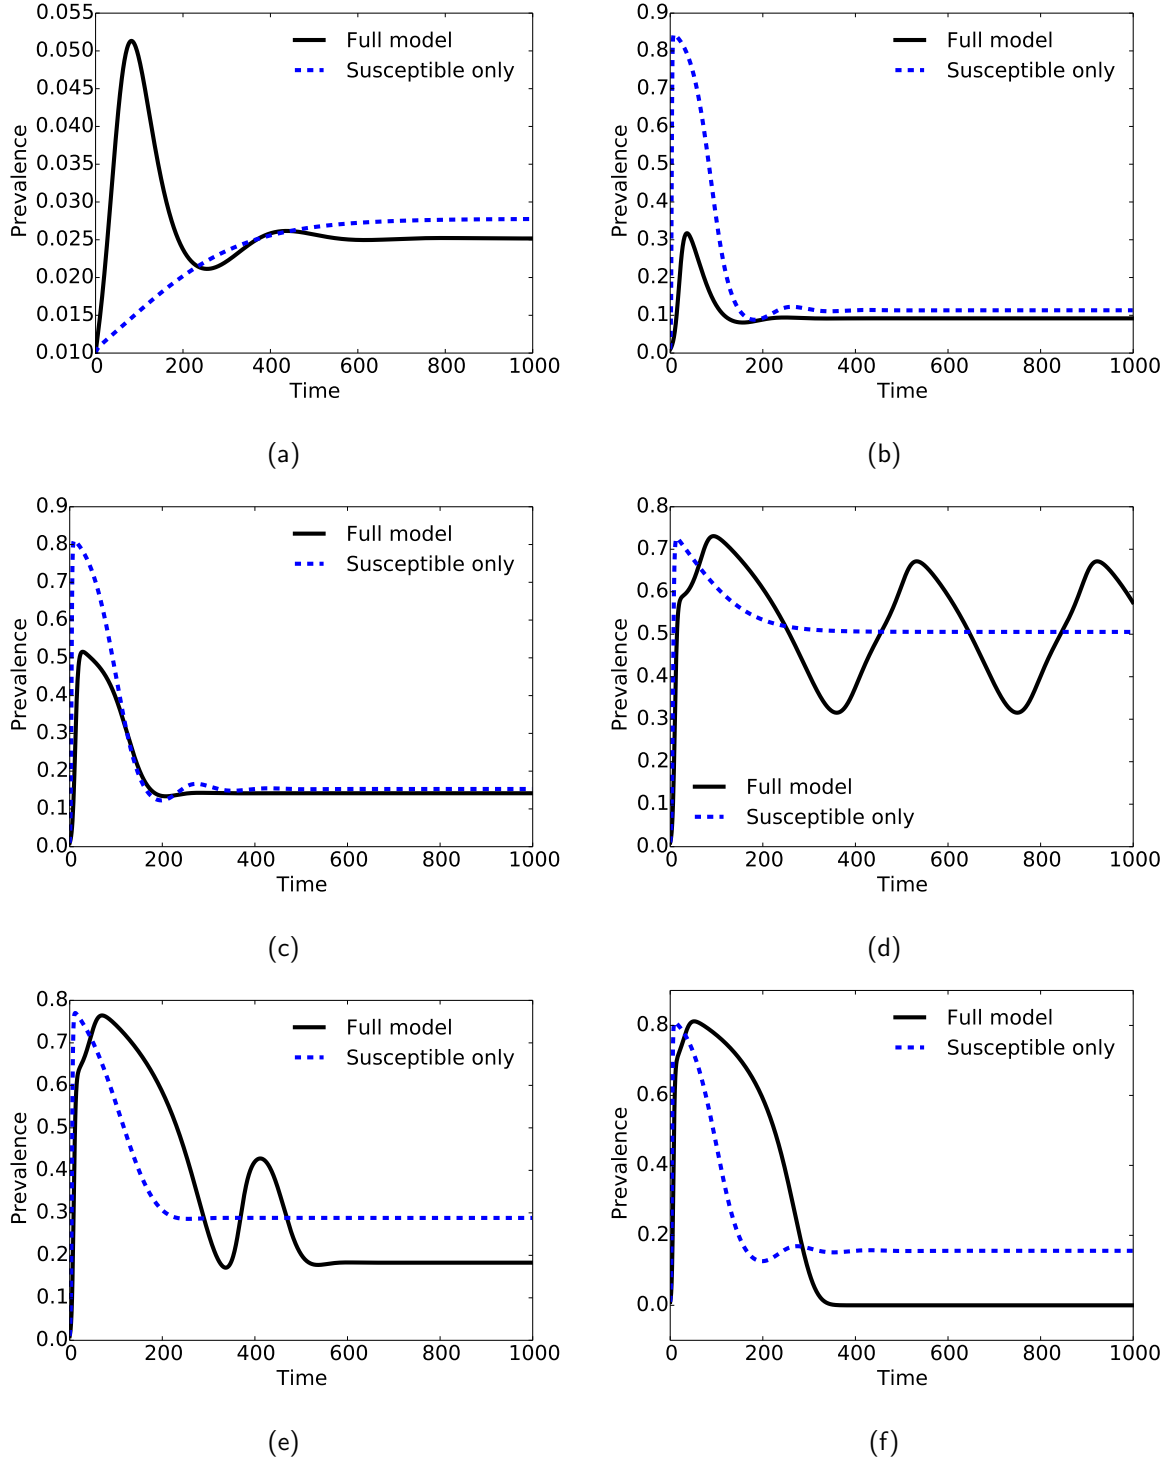


Figure D.11: The reduced model of susceptible-only behavior change fit to simulated trajectories from the full model with $\gamma = 0.5$, (a) Full model $\beta_b = 1.2$, best fit $\beta_b^* = 0.542$, (b) full model $\beta_b = 1.8$, best fit $\beta_b^* = 3.41$, (c) full model $\beta_b = 2.4$, best fit $\beta_b^* = 2.84$, (d) full model $\beta_b = 2.7$, best fit $\beta_b^* = 1.99$, (e) full model $\beta_b = 3$, best fit $\beta_b^* = 2.36$, (f) full model $\beta_b = 3.6$, best fit $\beta_b^* = 2.82$.




Figure E.12: Long term dynamics of the combined model where infected individuals always prefer unprotected sex for increasing values of \mathcal{R}_0 and the behavioral scale parameter s . (a) The amplitude of steady state prevalence oscillations. (b) The average prevalence at steady state.